

Method and Apparatus for Providing Information about a Real-World Space

Field of the Invention

- 5 The present invention relates to a method and apparatus for providing information about a real-world space by using virtual markers deposited in respect of users of the space.

Background of the Invention

- Ant colony optimization is concerned with the development and use of optimization
10 algorithms inspired by the collective behaviour of large colonies of social insects. Typically, the algorithms represent a problem space as a network of nodes connected by arcs. The network is traversed by simple, autonomous agents that are capable both of depositing virtual markers on nodes and arcs, and acting upon the accumulated markers that they encounter on their travels. By an appropriate choice of agent behaviours, some
15 emergent characteristic of the network can be caused to converge on a (near-)optimal solution to the original problem. This approach has been successfully applied to a range of problems such as the 'Traveling Salesman' Problem and job scheduling. More detailed information about the approach can be found in *Swarm Intelligence* by Bonabeau, Dorigo & Teraulaz (Oxford University Press, 1999).

- 20 Artificial life seeks to understand the processes by which biological and social complexity arise from simple organisms or agents. Typically, the approach is to construct a computer simulation of a universe populated by such agents and to study their interaction and evolution. The simulated universes may or may not reflect natural laws, and the agents may
25 or may not be modeled on naturally occurring organisms. Within that context, the simulation of ant-like agents with the capability to deposit and sense virtual markers (pheromones) has been known for at least a decade (for example, see *Ant Farm: Towards Simulated Evolution* by Collins & Jefferson (in *Artificial Life II*, Farmer et al, Addison Wesley, 1991).

- 30 Agent-based robotics applies similar ideas to motivate the development and exploration of swarms of simple interacting robots operating in the real world. The idea of pheromone

deposition and detection is well known in this field but is primarily used metaphorically to inspire mechanisms that actually implement *direct* communication between individuals rather than *indirect* communication through the environment in which the individuals move. For example, see *Progress in Pheromone Robotics* by Payton, Estkowski & Howard (preprint, 7th International Conference on Intelligent Autonomous Systems, March 25-27, 2002). An exception is the work of Andrew Russell in which robots do deposit (and sense) a chemical marker directly into the environment (see <http://www.ecse.monash.edu.au/staff/rar/>).

It is an object of a first aspect of the present invention to use the concept of pheromones to provide information concerning use of a real-world space such as an exhibition space.

Summary of the Invention

According to one aspect of the present invention, there is provided a method of providing information about a real-world space, comprising the steps of :

- (a) as the or each of at least one user moves through said space, virtual markers are deposited and stored to indicate associated locations visited by the user in the space, the virtual markers each having an initial strength value;
- (b) the strength values associated with the stored markers, either taken in location-dependent aggregations or individually, are caused to decay with time;
- (c) data about the current strength of the stored markers, either in aggregation or individually, is used to provide information relevant to use of the space.

Advantageously, the virtual markers deposited in respect of each user are deposited by a mobile device carried by the user; in this case, the virtual markers can conveniently be stored in a central system. Alternatively, the virtual markers can be arranged to be deposited and stored by an infrastructure system that monitors the locations of the users.

According to another aspect of the present invention, there is provided apparatus for providing information about a real-world space, the apparatus comprising:

- a first arrangement arranged to deposit and store virtual markers to indicate associated locations visited by each of multiple users in the space;

- a second arrangement arranged to decay with time the strength values associated with the stored markers, either taken in location-dependent aggregations or individually; and
- a third arrangement arranged to use data about the current strength of the stored markers, either in aggregation or individually, to provide information relevant to use of the space.

Brief Description of the Drawings

Embodiments of the invention will now be described, by way of non-limiting example, with reference to the accompanying diagrammatic drawings, in which:

- 10 . Figure 1 is a diagram of an exhibition hall having an arrangement for delivering relevant media objects to visitors in a timely manner as the visitors encounter items of interest in the hall;
- . Figure 2 is a diagram of a mobile device and service system used in the Figure 1 arrangement;
- 15 . Figure 3 is a diagram of a location report sent from the mobile device to the service system of Figure 2;
- . Figure 4 is a diagram of a response message sent by the service system to the mobile device of Figure 2;
- . Figure 5 is a diagram illustrating some of the choices available when implementing a pheromone trail mechanism such as provided by the mobile device and service system of Figure 2;
- 20 . Figure 6 is a diagram depicting the functional blocks involved in providing a pheromone trail mechanism;
- . Figure 7 is a diagram depicting a wake zone behind a user progressing around the Figure 1 hall;
- 25 . Figure 8 is a diagram illustrating the fall-off in allocated item usage probability with distance along and to the side of a movement track; and
- . Figure 9 is a table showing for each of multiple virtual features located around the Figure 1 exhibition hall, the likely next feature to be visited from the current feature based on counts of what previous users have done.
- 30

Best Mode of Carrying Out the Invention

Figure 1 depicts a real-world environment for which a number of zones have been defined in a virtual world that maps onto the environment. When a person moving in the environment (called a “user” below) is detected as moving into one of these zones, one or more media objects are delivered to the user via a communications infrastructure and a mobile device carried by the user. A zone may correspond to an area around a real-world object of interest with the media object(s) delivered to a user in this area relating to that real-world object. Alternatively, a zone may not correspond to any real-world object .

- 10 In considering such an arrangement, it is convenient, though not essential, to introduce the abstraction of a virtual feature which is the subject of each zone. Each such virtual feature is given a number of properties such as a unique identifier, a location in the real-world environment, the real-world extent of the zone associated with the feature, a subject description indicating what the feature concerns, and a set of one or more media-object
- 15 identifiers identifying the media objects (or “feature items”) associated with the feature. The zone associated with a virtual feature is referred to hereinafter as the ‘active zone’ of the feature.

For a feature that is intended to correspond to a particular real-world item (and typically having an active zone that maps to an area about a real-world object), this can be indicated in the subject description of the feature. Using the feature abstraction makes it easier to associate feature items that all relate to the same zone and also facilitates adding / removing these features items since data about the real-world extent of the related zone is kept with the feature and not each feature item.

25 Each feature is represented by a feature record held in a data-handling system, the feature records together defining the aforesaid virtual world that maps to the real-world environment. Each feature can be thought of as existing in this virtual world with some of these virtual features mapping to real-world objects.

30 As already noted, when a user is detected as within an active zone of a feature, one or more feature items are delivered to the mobile device of the user for presentation to the user. A

feature item can be presented automatically to the user upon delivery or the item can be cached and only presented upon the user having expressed an interest in the feature in some way such as by dwelling in the active zone of the feature more than a minimum time or by explicitly requesting presentation of the feature item. Indeed, the delivery of the feature item to the mobile device can also be deferred until the user is detected as having expressed an interest in the feature; however, since this approach can introduce a delay before the item is available for presentation, the embodiments described below deliver feature items to the mobile device of the user without awaiting a specific expression of interest in each feature (though, of course, a general filtering may be applied as to what items are delivered according what types of features are of interest to the user). Preferably, each feature or feature item is given a property indicating whether feature item delivery is to be effected automatically upon delivery or only after a user has expressed an interest in the feature; this enables important items (such as warning messages concerning features associated with potentially hazardous real-world items) to be pushed to the user whilst other items are subject to an expression of interest by the user. Advantageously, a user may elect to have feature items automatically presented even when the corresponding feature/item property does not require this. Furthermore, since as will be described hereinafter, pre-emptive caching of feature items in the user's mobile device may be implemented, automatic presentation is qualified so as only to apply where the user is in the active zone of the feature with which the feature item is associated.

Considering the Figure 1 example in more detail, the environment depicted is an exhibition hall 10 having rooms 11 to 17 where:

- room 11 is an entrance foyer with reception desk 18 but no associated virtual features;
- room 12 is a reference library with no associated virtual features;
- rooms 13, 14 and 15 are used for displaying real-world objects, namely paintings 20 and sculptures 21, for each of which there is a corresponding virtual feature centred on the object concerned and with an associated active zone 25 (indicated by a dashed line);
- room 16 is used for experiencing virtual features for which there are no corresponding real-world objects, the location associated with each feature being

indicated by a cross 22 and the corresponding active zone 25 by a dashed line; and

- room 17 is a cafeteria with no associated virtual features.

Virtual features are also defined in correspondence to the majority of openings 23 between rooms, the active zones 25 associated with the features again been indicated by dashed lines. Typically, the feature items associated with these features are incidental information concerning the room about to be entered and are automatically presented. It will be seen from Figure 1 that only a single feature is applied to an opening 23 so that it is not possible to tell simply from the fact that a user is detected in the active zone of the feature which room the user is about to enter; however, as will be later described, it is possible to determine from the user's past activity (either location based or feature based) the general direction of progression of the user and therefore which room is about to be entered. This enables the appropriate feature item to be selected for delivery to the user from amongst the items associated with the feature.

On entering the exhibition hall 10, a user 30 collects a mobile device 31 from the reception desk 18 (or the user may have their own device). This device 31 cooperates with location-related infrastructure to permit the location of the user in the hall 10 to be determined. A number of techniques exist for enabling the location of the user to be determined with reasonable accuracy and any such technique can be used; in the present example, the technique used is based on an array of ultrasonic emitters 33 (represented in Figure 1 by black triangles) positioned at known locations in each room (typically suspended above human level). The emitters 33 are controlled by controller 32 to send out emitter-specific emissions at timing reference points that are indicated to the mobile device 31 by a corresponding radio signal sent by the controller 32. The device 31 is capable of receiving both the timing reference signals and the emissions from the ultrasonic transmitters 33. The device 31 is also pre-programmed with the locations of these emitters and is therefore able to calculate its current location on the basis of the time of receipt of the emissions from the different emitters relative to the timing reference points.

The exhibition hall is equipped with a wireless LAN infrastructure 36 comprising a distribution system and access points 37. The wireless LAN has a coverage encompassing substantially all of the hall 10, the boundary of the coverage being indicated by chain-

dashed line 38 in Figure 1. The wireless LAN enables the mobile device to communicate with a service system 35 to download feature items appropriate to the feature (if any) corresponding to the current location of the user. In the present example, the determination of when the location of the user (as determined by the device in the manner already described) places the user within the active zone of a virtual feature, is effected by the service system; however, it is also possible to have the device 31 carry out this determination provided it is supplied with the appropriate information about the feature zones.

It will be appreciated that communication between the device 31 and service system 35 can be effected by any suitable means and is not limited to being a wireless LAN.

Figure 2 shows the mobile device 31 and service system 35 in more detail. More particularly, the mobile device 31 comprises the following functional blocks:

- A location determination subsystem 40 with an associated timing reference receiver 41 and ultrasonic receiver 42 for receiving the timing reference signals from the location infrastructure 32 and the emissions from the ultrasonic emitters 33 respectively; the location determination subsystem 40 is operative to use the outputs of the receivers 41 and 42 to determine the location of the mobile device (as already described above) and to send location reports to the service system 35.
- A visit data memory 43 for holding data about the current “visit” – that is, the current tour of the hall 10 being undertaken by the user of the mobile device 31.
- A feature-item cache 44 for caching feature items delivered to the mobile device 31 from the service system 35. The cache 44 has an associated cache manager 45.
- A communications interface 46 for enabling communication between the mobile device 31 and the service system 35 via the wireless LAN infrastructure 36.
- A user interface 48 which may be visual and/or sound based; in one preferred embodiment the output to the user is via stereo headphones 60.
- A visit manager 47 typically in the form of a software application for providing control and coordination of the other functions of the mobile device 31 in accordance with input from the user and the service system 35.
- A visit path guide 49 for giving the user instructions / indicators for following a

planned route around the hall 10.

Much of the foregoing functionality will typically be provided by a program-controlled general purpose processor though other implementations are, of course, possible.

- 5 The visit data held by memory 44 will typically include a user/device profile data (for example, indicating the subjects of interest to the user, the intended visit duration, and the media types that can be handled by the device), an electronic map of the hall 10, the user's current location as determined by the subsystem 40, and the identity of the feature (if any) currently being visited together with the IDs of its related feature items. The visit data also
- 10 includes a feature history for the visit, which is either:
- the history of visited features and their related feature item IDs in the order the features were visited (thus, a feature is added to the top of the visited-feature history list when the feature is encountered), or
 - the history of accessed features and their related feature item IDs in the order the
- 15 features were visited (thus, a feature is added to the top of the accessed-feature history list when one of its feature items is accessed by - that is, presented to - the user whilst the feature is the currently visited feature).

If a visited-feature history list is kept, a history of accessed features can be embedded in it by providing each feature in the history with an associated flag to indicate whether or not

20 the feature was accessed whilst current. Although keeping a visited-feature history provides more information about the visit, it will inevitably use more memory resources than an accessed-feature history and in many cases it will only be desired to track features which the user has found sufficiently of interest to access an associated feature item. Where the purpose of the feature history is simply to keep a list of features (and related feature

25 items) that were of interest to the user, it may be desirable to exclude from the list features for which items were automatically presented but are not associated with exhibits (real or virtual) - that is, exclude features concerned with incidental information about the hall.

The feature history preferably covers the whole of the visit though it may alternatively only

30 cover the most recently visited/accessed features. In either case, the most recent several entries in the history list form what is hereinafter referred to as the "feature tail" of the user and provides useful information about the path being taken by the user.

The visit data held in memory 43 may further include details of a planned route being followed by the user, and a history of the locations visited by the user (this may be a full history or just the locations most recently visited - hereinafter termed the “location tail” of the user).

The service system 35 comprises the following main functional elements:

- A communications interface 50 for communicating with the mobile device 50 via the wireless LAN infrastructure 36.
- 10 - An internal LAN 51 (or other interconnect arrangement) for interconnecting the functional elements of the service system.
- A data store 52 for storing feature data and, in particular, a feature record for each feature with each record comprising the feature identifier, the subject of the feature, the corresponding real-world location and extent of the feature’s active zone, the IDs and media type of the or each associated feature item, and a flag which when set
15 indicates that feature item presentation of an associated feature item is to be effected automatically upon delivery when the feature is being visited.
- A feature-item server 53 for serving an identified feature item to the mobile device 31 in response to a request from the latter.
- 20 - A location report manager 54 for receiving location reports from the location determination subsystem 40 of the mobile device and for passing on data from the reports to functional elements 55 and 56 (see below).
- A pheromone trail subsystem 55 for receiving location data, via manager 54, from all user mobile devices to build up trail data in a manner akin to the use of pheromones
25 by ants.
- An item-data response subsystem 56 for receiving location and other data from the manager 54 in order to prepare and send a response back to the mobile device 31 that provided the location data, about what feature items it needs, or is likely to need, both now, in view of a feature currently being visited, and (where, as in the present
30 embodiment, pre-emptive caching is implemented) in the near future. Subsystem 56 comprises a location-to-feature item translation unit 57 which can either be implemented independently of the data held in store 52 or, preferably, be arranged to

operate by querying the store 52, the latter having associated functionality for responding to such queries. Subsystem 56 further comprises a prediction unit 58 for predicting, in any of a variety of ways to be described hereinafter, what feature items are most likely to be needed in the near future.

- 5 - A route planner 59 for responding to requests from the mobile device 31 for a route to follow to meet certain constraints supplied by the user (such as topics of interest, time available, person or tour to follow, an exhibit or facility to be visited, etc). In providing a planned route, the route planner will typically access data from one or both of the feature data store 52 and the pheromone trail subsystem 55. The route
10 planner 59 can conveniently hold a master map of the hall 10 for use by itself and the other elements of the service system 35, and for download to each mobile device 31 at the start of each new visit and/or whenever the master map is changed.

The functional elements of the service system 35 can be configured as a set of servers all connected to the LAN 51 or be arranged in any other suitable manner as will be apparent to
15 persons skilled.

The mobile device 31 and service system 35 provide a number of useful capabilities that will each be described in detail below after an overview of the general operation of the mobile device and service system during a visit. It is to be understood that the split of
20 functionality between the mobile device 31 and service subsystem 35 can be varied substantially from that indicated for the Figure 2 embodiment; indeed all functionality can be provided either entirely by the mobile device 31 (with all feature items being stored in the device) or by the service system 35 (with the presentation of feature items to a user being by means of fixed input/output devices located around the hall near the locations
25 associated with the virtual features).

In general terms, a user starting a visit can request a route to follow using the user interface 48 of the mobile device 31 to indicate parameters to be satisfied by the route. This route request is sent by the visit manager to route planner 50 and results in the download to the
30 mobile device 31 of a planned route. The path guide 49 then provides the user (typically, though not necessarily, only when asked) with guide indications to assist the user in following the planned route . Where the interface 48 includes a visual display, this can

conveniently be done by displaying a map showing the user's current location and the planned route; in contrast, where only an audio interface is available, this can be done by audio cues to indicate the direction to follow. A user need not request a planned route and in this case will receive no guide indications. A user may request a route plan at any stage
 5 of a visit (for example a route to an exhibit of interest).

As the user moves through the hall, the location determination subsystem 40 sends periodic location reports 62 (see Figure 3) to the location report manager 54 of the service system 35 via the wireless LAN 36. In addition to the user's current location, these reports
 10 typically include a user identifier (and possibly, additionally or alternatively, a type identifier indicative of any variable of interest such as, for example, the group of users to which the device user belongs or an activity being undertaken by the user), user/device profile data, and prediction-assist data for use by the prediction unit 58 in predicting what feature items are likely to be needed shortly. This prediction-assist data can comprise one
 15 or more of the following: route data concerning any planned route being followed; the user's "location tail"; and the most recent feature (either the "most-recently visited" or "most-recently accessed") associated with the user, either provided alone or as part of the user's "feature tail".

20 When a location report 62 is received by the manager 54, it passes on the user's current location in the report to the pheromone trail subsystem 55 to enable the latter to build up trail data from all devices; additionally, the user and/or type identifier may be passed on to subsystem 55 if provided in the location report. The user's current location is also passed to the item-data response subsystem 56 together with any profile data and prediction-assist
 25 data in the location report 62. The item-data response subsystem 56 then constructs and sends a response 65 (see Figure 4) to the mobile device 31 that originated the location report.

More particularly, the location-item to feature translation unit 57 of subsystem 56 uses the
 30 data passed to subsystem to determine the feature, if any, currently being visited by the user and thus what feature items are relevant to the user in their current location. In doing this, the unit 57 may also use the supplied profile data to disregard both features that do not

relate to a subject of interest to the user and feature items of a media type that cannot be handled by the mobile device 31. The unit 57 may also use elements of the prediction-assist data (for example, the location or feature last encountered before the current one) to enable it to determine the direction of progression of the user and thus to select between
 5 feature items of a feature in dependence on the direction of approach of the user. This is done, for example, for the features associated with openings 25 in order to select a feature item appropriate to entering a room. The IDs of feature items identified by the unit 57 together with the identity of the corresponding feature and the status of the automatic presentation flag of the feature, form a first part 66 of the response 65 to be sent back to the
 10 mobile device 31. Where the current location does not correspond to the active zone of any feature, the first response part 66 simply indicates this.

A second part 67 of the item-data response 65 is produced by the prediction unit 58 and comprises a list of the feature items most likely to be needed in the near future by the
 15 mobile device 31; for each such feature item, the second response part 67 includes the feature ID, its type, size and probability of usage (discussed in detail hereinafter). Like the unit 57, the unit 58 uses supplied profile data to disregard feature items of features not of interest to the user or of a media type that cannot not be handled by the mobile device 31. The number of feature items identified in response part 67 is preferably limited (for
 20 example, to ten such items). The item-data response subsystem 56 then sends the response 65 back to the mobile device 31 of the user by using a return address supplied with the original location report 62 and passed to subsystem 56 by the manager 54.

Rather than having the prediction unit 58 provide a prediction each and every time the
 25 mobile device 31 sends a location report, it is possible to arrange for the prediction unit 58 only to operate when required by the mobile device 31 with the latter only requiring a prediction, for example, every nth location report or only after the user has moved a certain distance since the last prediction made by unit 58. Conveniently, the location report field used to carry the prediction-assist data is also used to indicate when a prediction is required
 30 by, for example, setting the field to a predetermined value when prediction is not required.

The item-data response received back at the mobile device 31 is processed by the visit

manager 47. If the first part 66 of the response identifies a feature (thereby indicating that the current location of the user corresponds to the active zone of feature), the manager 47 updates the 'current feature' data in memory 45 to the feature identifier and item IDs in the first response part. These item IDs are also passed to the cache manager 45 and are used by the latter to request immediate delivery of these items from the server 53 of the service system to cache 44, if not already present in the cache. If the feature history data held by memory 43 relates to visited, rather than accessed, features, and if the feature identifier and item IDs in the first response part 66 differ from the most recent entry in the feature history list, the latter is updated with the feature identifier and item IDs from the first response part 66.

In the case that no feature is identified in the first part of the response 65, the 'current feature' data in memory 43 is set to null.

The manager 47 also determines whether the (first) feature item (if any) identified in the first response part 66 is to be immediately presented to the user, this determination taking account of the setting of the automatic presentation flag in the first part of the response, any user indication (stored, for example in the profile data) that all items are to be automatically presented, and any monitored indications of the user's interest in the currently-visited feature. Where a feature item identified in the first response part is to be immediately presented to the user, the manager 47 requests the item from the cache manager 45 (there may be a delay in the delivery of the item if it has not yet been received from the server 53). At the same time, if the feature history concerns accessed features the manager 47 updates the feature history with an entry corresponding to the feature identifier and item IDs forming the 'current feature' data; where the feature history although recording all visited features, provides for indicating whether a feature has been accessed, the manager updates the feature history accordingly.

With respect to the data contained in the second part 67 of the response 65, the visit manager simply passes this data to the cache manager 45 which determines whether or not to request from server 53 any of the items identified that are not already in the cache 44. The cache manager 47 in making this determination takes account of the probability that an

item will be needed in the near future and the available cache space. The cache manager 45 may decide to create additional cache space by flushing one or more items from the cache and/or by reducing the space they occupy, as will be more fully described hereinafter.

5

In this manner, the cache manager 45 seeks to ensure that the next item requested by the visit manager 47 as the user progresses to the next feature will already be in the cache 44.

Following the processing of an item-data response by the visit manager, whenever a feature
 10 item is accessed (presented) either as a result of the visit manager 47 determining that the current feature is of interest to the user or as result of the user specifically requesting the item (for example, after inspecting the list of items associated with the current feature), then where the feature history data records accessed feature information, the visit manager 47 checks if the feature associated with the accessed item is the current feature and, if so,
 15 updates the feature history to record the feature as an accessed one if not already so indicated.

The visit manager can also be arranged to keep a list in memory 43 of the individual feature items accessed.

20

Having described the general operation of the mobile device 31 and service system 35, a more detailed description will now be given of some of the functionality embodied in the arrangement of Figures 1 and 2.

25 Pheromone Trails

The location reports provided by the mobile device 31 and passed to the pheromone trail subsystem 55 serve as virtual markers in the virtual world corresponding to the hall environment. These markers are stored by the subsystem 55 and used to build up trail and other useful information about utilisation of the corresponding real-world environment.

30

In general terms (that is, without limitation to the specifics of the embodiment of Figures 1 and 2), the virtual markers left in whatever manner in respect of a user can be given a

variety of characteristics. For example, the markers can be arranged to reflect the nature of pheromones laid by social insects such as ants and have the following characteristics:

- the markers are left automatically;
- markers from different users are undifferentiated;
- 5 - markers have a value that diminishes both with time and with the distance from the point of marking;
- markers accumulate, that is the value of overlapping markers at a point is the sum of their values at that point;
- markers can be detected by all other users of mobile devices in the environment.

10

However, each of these characteristics represents a choice in some dimension and other choices are possible. For example:

- each marker may be left following a specific user action to do so in respect of that marker (that is, left deliberately);
- 15 - markers may be identified by their source;
- markers may be of different types to reflect different activities or intentions by the source;
- markers may be persistent;
- markers may be stored as distinct elements;
- 20 - perception of the markers may be limited to particular users.

Of course, a wide range of mixes of the above choices of characteristics (and of other characteristics) are possible and although the term “pheromone trail” is used herein to refer to the general arrangement of the deposition and use of virtual markers, this term should not be taken as implying that any particular characteristic is being implemented in respect of the arrangement concerned or that the use of the markers is related to delimiting a trail . Furthermore, it is to be understood that implementation of any particular characteristic is not tied to either one of the mobile device 31 or service system 35. Indeed, the service system is not essential for the implementation of a pheromone trail arrangement where the devices can communicate amongst themselves. Conversely, whilst some form of mobile device will generally need to be carried by the user to assist in determining the location of a user, the actual location determination of a user and corresponding marker deposition can

- be done by the service system 35; for example, the user's mobile device can be arranged to emit distinctive ultrasonic signals at intervals which are picked up by fixed receivers with time of receipt information then being used to determine the user's location and a corresponding virtual marker deposited in respect of the user. A system that does not
- 5 require any device to be carried by the user for the purposes of location determination is also possible such as a camera-based system that can track the movement of an individual user through the hall 10 with the images produced by different cameras being correlated to follow the user as he/she passes from the field of view of one camera to that of another (this correlation can be aided by the use of face recognition technology). An alternative
- 10 approach is to use pressure sensors to detect the footfalls of users with the individual footfalls being correlated to determine the most likely pattern of related footfalls corresponding to movement of single users (this correlation is facilitated if the pressure sensors also give a weight reading for each footfall).
- 15 Whatever the detailed characteristics of the markers, the effect of their deposition as users move around the physical environment is the generation of a marker "landscape" in the digital representation of that environment. The ridges, peaks, troughs and wells of this landscape reflect the number of markers laid in each part of the landscape and will typically (though not necessarily in all cases) also reflect the time elapsed since the markers
- 20 were laid. The devices of other users are arranged to be able to sense this landscape enabling them to use various gradient and contour following applications to traverse it, for example to follow (or avoid) popular paths. In doing so, the intensity of marker accumulations can be indicated to users in a variety of ways; for example intensity levels can be represented through an audio signal whose loudness or frequency varied with that
- 25 intensity or through a visual display.

Figure 5 depicts some of the implementation choices available when constructing an embodiment of the pheromone trail arrangement, these choices being arranged by processing stage according to a division of the pheromone trail process into five such

30 stages (other divisions being possible). The five stages depicted in Figure 5 are marker deposition 80, storage 81, intrinsic behaviour 82 (applied to the stored data), application processing 83, and presentation. 84. These stages are carried out by corresponding

functional blocks 85 to 89 depicted in Figure 6 with the storage block 86 having two sub blocks, namely a storage pre-processing block 90 and a memory block 91. Also shown, in dashed lines, in Figure 6 are the mobile device 31 and pheromone trail subsystem 55 of the Figure 2 embodiment positioned to indicate where the functional blocks 85 to 89 are disposed in that embodiment.

Considering first the marker deposition stage 80 (functional block 85), marker deposition can be done automatically, by deliberate user-action per marker, or by deliberate user confirmation of an automatically-generated series of latent markers representing a trail segment. Where markers are deposited (or generated) automatically, the frequency of deposition/ generation can be made time or distance dependent (with respect to the last marker) or can be arranged to occur at specific way points around the environment, for example, at virtual features (that is, when a user enters the active zone of the feature, with typically only one marker being deposited/generated per feature encounter). Depositing a marker when a feature is encountered does, of course, require the mapping between location and features to have first been carried out; this can be done either by arranging for this mapping to be effected in the user's mobile device or by arranging for the unit carrying out the mapping away from the device (for example, unit 57 in the Figure 2 embodiment) to deposit a marker on behalf of the device.

However a marker is deposited/generated, each marker may have an associated user identifier and/or type indicator (indicating some special significance not specific to a user). In the case of there being more than one type of marker, either a single marker type can be associated with a user or multiple types of marker can be associated with the user. Where multiple marker types are associated with a user, different conditions can be specified for when each type of marker is to be deposited (for example, one type of marker could be deposited at regular intervals whilst another type only deposited when at a virtual feature). More than one type of marker can be deposited at the same time where appropriate and in this case it can be useful to avoid duplication of data by combining the different types of basic marker into a single compound marker with attributes defining the types of basic marker represented by the compound marker.

Each marker may also have a tag to indicate a desired decay behaviour – for example, where, by default, a marker is arranged to decay, a no-decay tag can be included which if set (or “true”) indicates that the marker concerned is not to be given the default behaviour of decaying in value with time. Of course, it is possible to view the decay tag simply as a marker type indicator with markers tagged for decay being decay-type markers and markers tagged not to decay being no-decay type markers.

The main choice presented at the storage stage 81 (functional block 86) is whether a marker is to be aggregated with previously stored markers deposited at the same location or stored as an individual marker along with any associated data. Whilst aggregation produces useful information, storing in an un-aggregated form has the advantage of preserving the maximum amount of raw data whilst leaving open the option to later on retrieve a copy of the data for aggregation; the disadvantages of not aggregating initially are the much greater storage capacity required and the delay in later on obtaining aggregated data. Where aggregation is effected, this can be done across all types of marker or for each type of marker separately. Typically aggregation is done by adding an initial strength value to the aggregated strength value already stored for the same “location cell” as that in which the marker was deposited where a location cell corresponds to a specific area of the real-world environment. Rather than a marker being allocated to one location cell only, the strength of the marker can be divided up between the nearest cells in proportion, for example, to the distance between the point of deposition of the marker and a center point of each of the nearest cells. The initial strength value of a marker can be made dependent on the type of marker concerned where multiple marker types are present.

The intrinsic behaviour stage 82 (functional block 87) applies behaviours to the aggregated or non-aggregated markers. For example, the aggregated or non-aggregated marker strength can be reduced with time with the rate of strength decay being dependent on marker type (the decay rate can be of any suitable form such as by a fixed amount per unit time or a by fixed proportion of the remaining strength per unit time). Where a marker is individually stored, the marker can also be given a limited life determined as expired either when its strength falls below a certain level or when the marker reaches a certain age (for which purpose, a time-of-deposition time stamp can be stored along with the marker). Applying

intrinsic behaviour is done, for example, by a process that regularly scans the memory block 91, reviewing and modifying its contents as necessary. The intrinsic behaviour stage 82 may not be present either because no relevant behaviours are being implemented or because they are applied as part of the applications processes for using the stored data.

5

The application stage 83 (functional block 88) uses the stored data to carry out specific application tasks and may also apply behaviours, such as marker strength fall off with time, to the data copied from storage where such behaviours have not been applied earlier to the stored data itself. Typical application tasks include:

- 10 - where markers of one or more types are aggregated (either on storage or by the application), determining and following a “ridge” of the highest-strength marker aggregations corresponding to the most popular trail through the environment; a related application is one where a “trough” of the weakest (or zero) marker aggregations is followed;
- 15 - where markers are stored individually with user IDs and a strength fall-off with time behaviour has been applied to the stored data, following a trail left by a specific user, the strength of the relevant markers indicating the direction of movement along the trail;
- where markers are stored individually with user IDs and deposition timestamps
- 20 enabling the trail laid down by each user to be followed, predicting or proposing a user’s future movement on the basis of the movement forward from that user’s current location of previous users whose trail leading to this location matches closely with the location tail of the subject user (that is, with the locations of the last several markers deposited by the current user);
- 25 - where markers are deposited on encountering a virtual feature and the markers are aggregated by type with a decay that is exponential in form with a time constant of half a day for example, determining the most popular features of a given type for the current day by determining the strongest aggregation of markers of that given type.

It should be noted that different applications may call for different marker strength decay

30 rates. This can be accommodated in a several ways – for example, each marker that is deposited can be split into multiple copies with each copy being used for a particular application and decayed (either as an intrinsic behaviour or by the application) at an

appropriate rate. A variant of this approach is to give a single marker multiple strength attribute values, each value being associated with a different application and being decayed at a rate appropriate for the application concerned either as an intrinsic behaviour or by the application; this is equivalent to there being a respective marker type per application with
 5 markers of several different types being deposited at the same time in a compound marker (of course, it would also be possible to actually deposit discrete markers per application type).

As regards the presentation stage 84 (functional block 89), how the output of an application
 10 is presented to a user will depend on the nature of that output and the interface modalities available. Typically, where an application task has determined a trail to follow or the most popular features, this can be presented to the user on a map display whilst if an application is arranged to provide real time guidance along a path, this may be best done using audio prompts.

15 Implementation details of the functional blocks 85-89 for any particular combination of the available choices discussed above will be apparent to persons skilled in the art. It should be noted that multiple combinations of choices can exist together; for example, markers can be arranged to be deposited by a mobile device both at fixed time intervals and every time
 20 a feature is encountered and a marker can be both aggregated upon storage as well as an individual copy being kept. Thus in one implementation, an array data structure is used to define an X-Y array of location cells with each cell mapping to a respective area of the real world environment (hall 10) and being used to hold, for each marker type, both an aggregated strength value for the markers of that type deposited at locations covered by the
 25 real-world area to which the cell maps, and a pointer to a linked list of individual records for those markers which are still alive (that is, not time expired).

With respect to the embodiment of Figures 1 and 2, the pheromone trail subsystem 55 is
 30 arranged to store markers of three different types, namely:

- **“tour” markers** deposited in the form of location reports 62 by a tour guide and serving to delineate a proposed route around the hall. These markers are each

deposited by deliberate act of the tour guide and have an associated “no-decay” tag as well as an indicator of their type. Preferably the type indicator has an associated sub-type that identifies a specific tour. Since each specific tour will have only one set of markers associated with it, storing the tour markers on the basis of aggregating markers of the same type and sub-type deposited in the same location is the same as storing the markers individually and either approach may be adopted. The stored markers are not decayed with time.

- **“normal” markers** deposited in the form of location reports 62 by the mobile devices 31 of visitors, these markers being deposited at fixed time intervals and being subject to aggregation on storage with other markers of this type deposited in the same location cell (that is, an initial strength value associated with a newly deposited marker is added to the aggregated strength value associated with the marker aggregation for the cell in which the new marker has been deposited). The strengths of the marker aggregations are decayed with time but over a long time period. These aggregated “normal” markers serve to indicate the most popular trails, reflecting both the number of users traversing these trails and the time spent on them.
- **“feature” markers** deposited by the unit 57 each time it determines from data in a location report that the device sending the report is in the active zone of a feature. If, as is preferred, the prediction assist data in the location report contains current feature data, then deposition of a feature marker can be restricted to when a user first enters the active zone of the feature, this being achieved by comparing the identity of the current feature as determined by unit 57 with the current feature noted in the location report and only depositing a marker if the two differ. The feature markers are aggregated in feature cells held by the unit 55 and are decayed over a period of an hour to give a picture of the current popularity of the features. Feature cells are simply location cells covering an area corresponding to the active zone of a feature. In a variant, the deposition of a feature marker is only effected when a user is in the active zone of a feature and requests presentation of a related feature item.

The stored markers are put to use for route planning / following, feature popularity review, and prediction purposes. With respect to route planning, when the visit manager 47 of a mobile device 31 requests a route from the route planner 59 of the service system, the latter

can ask the application task block 88 of the pheromone trail subsystem 55 to access the stored marker data and propose a possible route based either on the tour markers or the aggregated normal markers. Thus, the route planner, where provided with a subject of interest to the user by the visit manager 47, can be arranged to map this subject to a particular tour sub-type and then retrieve the set of locations of the corresponding tour markers stored by the subsystem 55; these locations are then used to provide a route plan back to the mobile device 31. As described above, no sequence information is stored with the tour markers and whilst this will generally not be an issue, it is possible to provide for the tour markers to carry sequence information in a number of ways, the simplest of which is to associate a sequence number with each tour marker as it is deposited, this number being incremented for each successive marker and being stored along with the marker. An equivalent way of providing sequence information is to incrementally increase/decrease the strength value assigned to each marker as it is deposited; since the tour marker do not decay, this strength value remains and effectively serves as a sequence number indicating the direction of progression of the tour.

The route planner 59 can be arranged to request the subsystem 55 for the most popular route around the hall 10 as indicated by ridges of higher-strength accumulations of normal markers, or for the least crowded route as indicated by troughs of zero or low-strength accumulations of the normal markers. Of course, the route planner 59 will typically have been requested by a user to provide a route that takes the user to features relating to a particular subject or even to a set of user-selected features; if the route planner decides that there is no relevant pre-planned tour it can use, or if the user has specifically asked for a popular or a least crowded route, then the route planner will use the normal-marker aggregations to aid it in planning a route between the selected features. This can be done by first selecting an order in which to visit the features and then asking the application task block 88 to provide the most popular / least crowded route between each successive pairing of features in the order they are to be visited. Alternatively, the actual order of visiting of the features, as well as the route between each feature, can be determined according to the peaks and troughs of the accumulated normal marker landscape, preferably with account being taken of the total distance to be traveled by the user. In this case, since the application task block 88 has more immediate access to the stored marker accumulations, it

may be appropriate for the route planner to hand over the whole task of planning a route to the task block 88.

Rather than determining a route by following ridges or troughs in the accumulated-marker landscape, the route planner can be arranged to determine a route by avoiding ridges or troughs. In relation to route determination, it is to be understood that the term “ridge” includes the limit case of a “peak” and the term “trough” includes the limit case of a “well”.

- 10 An image of the virtual landscape formed by the location-dependent aggregations of markers mapped to a representation (such as a plan) of the hall 10 can, of course, be passed to the mobile device 31 for presentation to the user.

- 15 Another possible usage of the pheromone trail subsystem 55 in respect of providing route information involves the deposition by a first user of user-specific markers that are not aggregated but are arranged to decay in strength over a period of an hour or so. These markers would enable a second user to request the route taken by the first user (for example, by means of a request sent from the visit manager 47 of the second user’s mobile device to the route planner 59), the markers deposited by the first user then being accessed
20 to determine the route taken by the first user and their direction of progression as indicated by the current strengths of the markers. This service (suitable for a parent wanting to track a child) can be made private with only the users involved being able to access the relevant marker data and can be provided as a fee-based service.

- 25 A similar type of usage involves all members of a group having markers of a type specific to that group, the markers being aggregated on storage. This would enable an overview to be obtained of what the group did during a visit and if the markers concerned did not decay (though typically given a lifespan limited to the day of the visit) and were deposited at fixed time intervals, it would also enable the popularity of different visited features to be
30 discerned. Preferably, the group markers are deposited in addition to normal markers rather than as an alternative to the latter.

Although in the foregoing examples of the use of the pheromone trail system in the embodiment of Figures 1 and 2, the route information derived from the stored markers has been passed back to the mobile device for storage in the visit data memory 43 as a route to be followed, it is also possible to have a more dynamic interaction between the mobile device and the stored marker data. Thus, for example, the mobile device 31 can be arranged to query the pheromone trail subsystem 55 continually as to the next location to move to in order to follow a ridge or trough of the marker landscape or to follow a trail laid down by a specific user.

With regard to the use of the deposited marker data for feature popularity review, if a user wishes to ascertain the current relative popularity of the features (or, in user terms, of the exhibits with which the features are associated), the user causes the visit manager 47 to send a request to the pheromone trail subsystem 55. The task block 88 of the subsystem 55 then accesses the feature marker accumulations of the feature cells and uses the strengths of these accumulations to determine the current relative popularity of the features. This popularity data is then returned to the requesting mobile device for presentation to the user. If a longer term view of the popularity of the features is required, then the task block 88 accesses the normal marker aggregations for the feature cells, these aggregations having a longer decay period and, unlike the feature marker accumulations, having a strength that reflects the dwell time at each feature as well as the number of visits.

In respect of use of the deposited marker data for prediction purposes, this involves using the current location or location tail of a user to make predictions as to where the user is likely to go next having regard to what others have done as indicated by the relative strengths of the accumulations of normal markers in location cells adjacent the one in which the user is currently located. If location tail data is available, the strengths of marker accumulations in location cells just visited by the user (and possibly also of the cells on either side of such cells) can be scaled down to reflect the fact that the user is less likely to visit those cells; however, if the geography of the hall or the layout of features of interest to the user is likely to cause the user to turn around, then such scaling down is not effected. Making predictions of the user's future path in this manner is carried out by the application task block 88 of the pheromone trail subsystem. As will be further described below, this

future path prediction capability can be used by the prediction unit 58 to determine what feature items are likely to be needed in the near future.

It will be appreciated that many other applications are possible for the pheromone trail
5 arrangements discussed above.

With respect to management of the pheromone trail information by the exhibition hall staff, the use of tour markers for defining tours has already been mentioned. However, other management techniques are also possible. For example, as an alternative to using tour
10 markers, or in order to modify trails such as those defined by aggregation ridges, a special marker type that has a very high initial strength can be defined and associated with a tour guide – this guide then traverses the hall on a desired path depositing the high-strength markers along the way. These high-strength markers effectively serve to swamp existing trail information to define new trails. A reverse effect can also be provided by defining a
15 negative strength marker type to wipe out (or at least reduce) aggregated strength values along particular paths.

Caching of Feature Items

As described above, the mobile device 31 is arranged to pre-emptively cache feature items
20 in cache 44 in dependence on their respective probabilities of being required in the near future; these probabilities being determined by the prediction unit 58 of the service system 35 using information (in particular, the prediction-assist data) provided in the location reports from the mobile device 31. In this manner, the latency inherent in fetching feature items from the feature item server 53 only when needed is avoided.

25

The prediction unit 58 can operate on the basis of any one or more of a variety of different techniques for predicting which feature items will be needed in the near future. A number of these techniques are described below, these techniques being divided into two groups, namely a first group A covering techniques that do not use visit data concerning previous
30 users, and a second group B that rely on such previous-users visit data.

It should be noted that in the following the probability of an item being needed (also

referred to as the probability of usage of the item), is used to encompass both the probability of an item being definitively requested (that is, not on a probabilistic basis) for delivery to the mobile device and the probability of an item being accessed at the device by the user. The fact that these two probabilities are different in the Figure 2 embodiment is

5 because the service system and mobile device operate on the basis that all items associated with a currently visited feature are downloaded into the cache 44 of the device, regardless of their probability of being accessed by the user. The probability that a particular item will be requested for delivery to the device is thus the same as the probability that the user will visit the associated feature. Had the service system and mobile device simply been

10 arranged to non-probabilistically request delivery of an item only when accessed by the user, the probability of an item being requested for delivery would be the same as the probability of that item being accessed. Notwithstanding the fact that in the Figure 2 embodiment all items associated with a current feature item are requested for delivery, prediction of what items may be needed in the near future need not be restricted to use of

15 the probability of a feature item being non-probabilistically requested (as indicated, for example, by the probability of the associated feature being visited), and can alternatively be based on the probability of the user accessing a particular item (or of accessing at least one of the items associated with a feature, all these items then being considered as having the same probability of access). Consolidating the foregoing, the probability of usage of an

20 item can be based on the probability of a feature being visited or accessed, or of a feature item being accessed by the user or non-probabilistically requested for delivery.

Furthermore, regardless of the prediction technique being used, the prediction unit 58 may, as already mentioned, filter out from its prediction process all feature items that do not

25 relate to a subject of interest to the current user or are of a media type incompatible with the mobile device of that user. In fact, rather than filtering out all feature items concerning subjects in which the user has not expressed interest, the probabilities associated with these items regarding their likely use in the near future can be appropriately adjusted to take account of the user's apparent lack of interest in them.

30

A – Prediction not based on Visit Data Concerning Previous Users

1. ADJACENCY OF FEATURES TO CURRENT LOCATION (OPTIONALLY WEIGHTED AGAINST WAKE FEATURES)

This is the simplest prediction technique and in its basic form takes the current location of the user and determines the closest features (typically by reference to the data held in the feature data store). The probability of usage of the feature items associated with these features is based on the probability of the features being visited and is thus set to fall off in dependence on the distance of the feature concerned from the user's current location. For example, all features within a 30 meter radius of the user's current location are determined and the probability of usage of an item associated with a feature r meters away is set to:

$$(30-r)/30$$

This basic technique can be modified to reduce the probabilities of usage of feature items associated with features that the user has recently passed (and is therefore less likely to visit in the immediate future). These features, referred to below as "wake" features, are identified by the prediction unit 58 using location history data of the current user – in present embodiment this data is supplied in the form of the user's location tail provided as part of the prediction assist data. As depicted in Figure 7, the immediately preceding location 130 (or locations) of the user are used, together with the user's current location 131 to determine a "wake" zone 133; the probability of usage of any feature item associated with a feature 134 lying in the wake zone 133 (that is, a wake feature) is then weighted by a factor between 0 and 1. It will be appreciated that the wake zone 133 could be divided into sub-zones each having a different associated weighting factor according to a perceived reduced probability of usage of feature items for features in such sub-zones.

25

2. ADJACENCY OF FEATURES TO PLANNED ROUTE

If the user is following a planned route and data about the next portion of this route is included in the prediction assist data, the prediction unit 58 can use the route forward of the user's current location to determine the features next to be encountered along the route (either on the route or adjacent to it). The probability P of usage of a feature item associated with such features is again based on the probability of a feature being visited and is set according to both the distance l of the feature along the forward

30

route (the greater the distance, the lower the probability) and the perpendicular distance d of the feature off the route (again, the greater the distance the lower the probability but this time the fall-off rate is much faster than for distance along the forward route). Figure 8 illustrate, by way of example, a linear fall off of probability with distances l and d giving a defining plane 139.

It may be noted that where the route being followed is a standard route for which route data is held by the route planner 59, then the route data included in the prediction assist data can simply be an identifier of the route, the prediction unit using this identifier to retrieve the route details from the route planner 59. It may also be noted that a planned route may be defined in terms of features to be visited rather than as a path; in this case, the probability of usage of feature items for features on the route is set simply by their order from the current point onwards; other features not on the route can still be included in the prediction according to their adjacency to the features on the route (or to a direct path between them). It may be further noted that having a planned route stored in visit memory 43 is not necessarily to be taken as a sufficient condition that the user is following a planned route; one or more additional conditions may be required such as, for example, the user is actively using the path guide unit 49 to follow the planned route, or the last two/three features visited have all been on the planned route. The determination as to whether a planned route is being followed is preferably made in the mobile device 31.

3. ADJACENCY OF FEATURES TO FUTURE TRACK PREDICTED FROM MOVEMENT HISTORY

Where the user's recent movement history is available to the prediction unit 58 (for example, as a result of the user's location tail being included in the prediction data), then the unit 58 can use this information to predict the user's track in the immediate future. Thus, if the user's location tail is available to the prediction unit 58, a smooth curve passing through the locations in this tail can be determined and continued to predict the user's future track. This track can then be used in much the same manner as a planned route as described above, that is, the features lying on or near the track are identified and the probability of usage of feature items associated with these features is set in dependence both on the distance of the features concerned along the

track and on their distance off the track (c.f. Figure 8).

Rather than predicting the user's future track on the basis of their location tail, this track can be predicted from a knowledge of the user's current location and direction of moving as determined for example, by the direction of facing of the user's body as measured by an electronic compass carried by the latter.

B – Prediction based on Visit Data Concerning Previous Users

This group of prediction techniques use visit data concerning previous users. This visit data can be collected in any suitable manner. For example, the visit data can be obtained by storing in a mobile device 31 during a visit, time-ordered lists of all locations and features visited, and all feature items accessed and where they were accessed. At the end of the visit, the stored data is uploaded to the service system for organization and use by the prediction unit 58. It is alternatively possible to arrange for the visit data to be collected by the service system as a user progresses through a visit. Furthermore, where prediction is based on location / feature trail information, the pheromone trail subsystem 55 can be used to provide the required visit data.

20 4(a). SAME LOCATION – TRACK PREDICTION - FEATURE PREDICTION (OPTIONALLY WEIGHTED AGAINST WAKE FEATURES)

This prediction technique simply uses the user's current location to predict where the user is likely to go next on the basis of where previous users have gone from this location (the prediction unit 58 may, for example, query the pheromone trail subsystem for such information). Given the most likely future track(s) of the user, the features that will be encountered along or near the track are determined followed by the probability of usage of the associated feature items; this is effected, for example, in a manner akin to that used in prediction technique (3) above. If more than one future track is considered, the probability of use of each track is used as an additional weighting factor for the probability of usage of the feature items. A weighting can also be introduced to reduce the probability of usage of feature items associated with wake features as described above with reference to prediction technique (1).

4(b). SAME LOCATION – DIRECT PREDICTION OF FEATURE/ FEATURE ITEM (OPTIONALLY WEIGHTED AGAINST WAKE FEATURES)

Rather than predicting feature/feature item usage indirectly by first predicting a future track for the user based on the tracks taken by previous user's from the current location, it is possible to take the user's current location and use it to predict directly from the previous-users visit data what features will probably be visited or accessed next - or even more directly, what feature items are likely to be visited / delivered to the mobile device / accessed in the near future. This is done by organizing the previous-users visit data on the basis of what features are most commonly next visited or accessed or what feature items are next delivered to or next accessed by users who have been at the current location. Again, a weighting can be introduced to reduce the probability of usage of feature items associated with wake features.

5(a). SAME RECENT MOVEMENT HISTORY - TRACK PREDICTION - FEATURE PREDICTION

This technique is similar to prediction technique (4a) but makes its future track prediction based on where previous users with the same recent movement history (typically, with the same location tail) have gone from the current location. Of course, since previous locations visited by the user are inherently taken into account by this technique, it is inappropriate to adversely weight the usage probabilities of items associated with wake features as was optionally done for prediction technique (4). It should be noted that in order to be able to identify previous users with the same recent movement history, the movement data of previous users needs to be available in an un-aggregated form.

5(b). SAME RECENT MOVEMENT HISTORY - DIRECT PREDICTION OF FEATURE/ FEATURE ITEM

This technique is similar to prediction technique (4b) but uses visit data from previous users with the same recent movement history (typically, with the same location tail) in order to determine what features are likely to be visited or accessed in the near future, or what feature items are likely to be visited / delivered to the mobile device / accessed in the near future.

6. SAME MOST-RECENTLY VISITED FEATURE - PREDICTION OF FEATURE/ FEATURE ITEM (OPTIONALLY WEIGHTED AGAINST RECENTLY-VISITED FEATURES)

This prediction technique does not use location data but bases itself on visited feature data. More particularly, the prediction unit 58 uses the current feature identified by the location-to-feature translation unit 57 or, if the user's current location does not correspond to a feature, the user's most-recently visited feature as identified in the prediction assist data included in the latest location report. Given this feature, the prediction unit 58 accesses a stored table 140 (see Figure 9) which for each feature $F1$ to F_N keeps a count of the next feature visited by each previous user. These count values enable the unit 58 to determine the probability of each feature being the next feature visited, these probabilities then being applied to the feature items associated with each feature as the probability of usage of those items. If the prediction assist data includes the feature tail of the user, this can be used to reduce the probabilities associated with the features in the user's feature tail.

The table 140 lends itself to dynamic updating since if the unit 57 identifies a feature - for example feature $F(N-3)$ - that is different to the most-recently visited feature - for example, feature $F5$ - identified in the prediction assist data, this indicates that the user has moved from the feature $F5$ to the feature $F(N-3)$ so that the count value in the table cell at the intersection of row $F5$ and column $F(N-3)$ should be incremented to reflect this.

It should be noted that a single access to the table 140 will only give probabilities regarding the next feature to be visited. However, it is possible to look further ahead by accessing the table again in respect of the most-probable next feature (or features) in order to derive probabilities in respect of the next-but-one feature to be visited. By repeating this process, a forward-looking probability graph can be built up to any required depth.

It should also be noted that it is possible to provide table 140 with the next-visited feature probability data replaced by next-accessed feature data where, as explained above, an accessed feature is a feature having at least one associated item that has

been accessed – presented to – the user. Alternatively, the next-visited feature count data in table 140 can be replaced by probability data about the next feature item visited / requested for delivery (or delivered) to / accessed by, the user after the current/most-recent feature visited.

5

7. SAME MOST-RECENTLY VISITED FEATURE HISTORY - PREDICTION OF FEATURE/ FEATURE ITEM

This prediction technique matches the user's recent visited-feature history (their visited-feature tail) to the visited-feature histories of previous users visiting the user's current or most-recently visited feature. Having identified previous users with matching feature tails, the prediction unit 58 analyses the visit data of these previous users to determine the probabilities associated with the user next visiting the other features and thus the probability of usage of the associated feature items. As with prediction technique (6), rather than predicting the next feature to be visited, the previous visit data can be organized to enable the unit 58 to predict the next accessed feature (and thus feature items likely to be needed) or the next feature item visited / requested for delivery (or delivered) / accessed.

10

15

20

8. SAME MOST-RECENTLY ACCESSED FEATURE - PREDICTION OF FEATURE/ FEATURE ITEM (OPTIONALLY WEIGHTED AGAINST ACCESSED WAKE FEATURES)

25

This prediction technique is similar to prediction technique (6) but is based on accessed features rather than visited features. It should be noted that since the location-to-feature translation unit 58 does not know whether any feature it identifies is an accessed feature, the prediction unit works on the basis of the most-recently accessed feature as identified in the prediction assist data it receives from the user.

30

Rather than predicting the next feature to be accessed, the previous visit data can be organized to enable the unit 58 to predict the next feature to be visited (and thus feature items likely to be needed) or the next feature item visited / requested for delivery (or delivered) / accessed.

9. SAME MOST-RECENTLY ACCESSED FEATURE HISTORY - PREDICTION OF FEATURE/
FEATURE ITEM

This prediction technique is similar to prediction technique (7) but is based on accessed features rather than visited features. Again, rather than predicting the next
5 feature to be accessed, the previous visit data can be organized to enable the unit 58 to predict the next feature to be visited (and thus feature items likely to be needed) or the next feature item visited / requested for delivery (or delivered) / accessed.

10 It will be appreciated that since each feature item can be represented by a respective feature, where the foregoing prediction techniques involve features the same techniques can be applied directly to feature items provided the latter have any required associated parameter data, such as location. Thus, the prediction techniques (1), (2), (3), (4a) and (5a)
15 which all involve determining the closeness of features to a location or track, can equally be implemented by determining the closeness of individual feature items to a location or track. Similarly, techniques (6) to (9) can be applied by accessing the previous-users visit histories in respect of the same visited/accessed feature item / feature-item tail rather than the same visited/accessed feature / feature tail (in this context, a “visited” feature item is one where the user has visited the location associated with the location).

20 Where an above-described prediction technique is based on determining the probability of visiting / accessing a feature, then instead of using this probability as the probability of usage of all the feature items concerned, it is alternatively possible to set the usage probability of each feature time individually by weighting the feature-related probability
25 according to the relative popularity (in terms of actual presentation to the user) of the item concerned with respect to other items associated with the same feature— provided, of course, that data about relative popularity is made available to the prediction unit 58.

All of the above prediction techniques can be implemented fully in service system 35, split
30 in any appropriate manner between the service system 35 and the mobile devices 31, or fully in the mobile devices 31, even if based on the visit histories of previous users. Thus, for example, where prediction is done on the basis of previous visit histories but there is no

service system 35, each mobile device can be arranged to store all its past visit histories and to supply them to other devices on request. As another example, the Figure 2 embodiment can be modified by arranging for the prediction unit 58 simply to provide the mobile device with the probabilities of features being visited / accessed, it then being up to
 5 the mobile device (in particular the cache manager 45) to translate features to feature items and request such items according to the probabilities associated with the corresponding features; this, of course requires the cache manager to have access to information about the association between the features and feature items and such information can conveniently be stored in memory 43. Rather than the cache manager 45 requesting individual items
 10 from the server 53 when effecting pre-emptive caching, it can supply a feature identifier to the server 53 which then returns all the feature items associated with the feature concerned.

Additional prediction techniques to those described above are also possible. Also, the above-described pre-emptive caching arrangement can equally be applied where the
 15 features items are being supplied to the cache 44 from a local storage device such as a DVD drive rather than from a remote resource over a wireless connection.

It is also possible to control loading of items into the cache 44 on the basis that they have not been identified as an item having a low probability of usage as determined using one of
 20 the above-described prediction techniques. In one implementation of this approach, the cache manager 45 is arranged by default to request from the server 53 all items associated with features within a predetermined distance of the user's current location (as determined, for example, by querying the feature data store 52); however, this default is overridden in respect of any item which, according to the prediction unit 52, has a probability of usage
 25 below a predetermined threshold value. In this example implementation, the prediction unit 52 is arranged to identify the low usage probability items based on the information received in the location reports 62 from the device 31, the identities of these items then being returned to the device 31 in a response message 65.

30 Other factors additional to item usage probability may be used to determine when an item should be loaded into cache. For example, the amount of free space in the cache can be used to control the threshold probability value below which items are not loaded into the

cache – the fuller the cache, the higher this threshold is set.

An item retrieved by the mobile device 31 to the cache 44 will typically be retained in cache for as long as possible to be available for access by the user at any time including
 5 after the user has passed on from the feature with which the item is associated. However, since the size of the cache memory 44 will generally be much smaller than that required to store all available feature items, it will usually be necessary to repeatedly remove items from the cache during the course of a visit to make room for other features items.

10 Items to be flushed from the cache can be identified on the basis of a prediction-based indication of what items are unlikely to be needed again. The above-described prediction techniques used for determining the probability of usage of feature items can also be used in determining whether a cached feature item should be flushed from the cache.

15 Variants

It will be appreciated that many variants are possible to the above described embodiments of the invention. For example, although in all the embodiments described above, all feature items have originated from the same source, namely, item server 53, it is also possible to provide for multiple item sources each holding a respective subset of the items. In this
 20 case, the item identifier associated with each item can be arranged to indicate directly the source from which the item can be obtained, or some other mechanism can be employed to direct an item request to the appropriate source. The multiple item sources effectively form a distributed item server.

25 As already noted, the distribution of functionality between mobile devices and the service system is not limited to the distributions described above since the availability of communication resources makes it possible to place functionality where most suitable from technical and commercial considerations. Furthermore, in the foregoing reference to a mobile device is not to be construed as requiring functionality housed in a single unit and
 30 the functionality associated with a mobile device can be provided by a local aggregation of units.

The above described methods and arrangements are not limited to use in exhibition halls or similar public or private buildings; the methods and arrangements disclosed can be applied not only to internal spaces but also to external spaces or combinations of the two.